MULTIMEDIA UNIVERSITY

STUDENT ID NO

MULTIMEDIA UNIVERSITY

FINAL EXAMINATION

TRIMESTER 3, 2016/2017

TSN2101 – OPERATING SYSTEMS

(All sections / Groups)

1 JUNE 2017
9.00 a.m - 11.00 a.m
(2 Hours)

INSTRUCTIONS TO STUDENTS

1. This Question paper consists of 8 pages with 7 Questions only.

2. Attempt **FIVE** out of **SEVEN** questions. All questions carry equal marks and the distribution of the marks for each question is given.

3. Please print all your answers in the Answer Booklet provided.

QUESTION 1

a) One way for the operating system (OS) to protect itself and other system components is to have dual-mode operation, user mode and kernel mode. Explain how having these two modes of operation function as a form of protection.

[3 marks]

b) For efficient performance and implementation an OS should be partitioned into separate subsystems. Discuss the layered approach to arranging these subsystems.

[4 marks]

c) Process scheduling keeps the CPU busy at all times and to deliver acceptable response times for all programs. Describe the differences between short-term, medium-term and long term schedulers.

[3 marks]

d) Depending on system implementation, a child process may receive some amount of shared resources with its parent. What are the two possibilities for the address space of the child relative to the parent?

[2 marks]

Continued ...

QUESTION 2

a) What resources are used when a thread is created? How do they differ from those used when a process is created?

[3 marks]

b) Consider the following set of processes:

Process	Arrival time	Burst-time	Priority
P1	0	10	3
P2	4	2	1
P3	3	3	3
P4	5	1	4
P5	1	4	2

i) Draw the Gantt charts illustrating the execution of these processes using First Come First Served, Non-Pre-emptive Priority (a smaller priority number implies a higher priority), Pre-emptive Shortest Job First (SJF) and Round Robin (RR) (quantum = 4) scheduling.

[1 + 1 + 2 + 2 = 6 marks]

ii) What is the turnaround time of P4, for the Non-Pre-emptive Priority, Pre-emptive SJF and RR algorithms?

[1.5 marks]

iii) What is the waiting time of P2, for the Non-Pre-emptive Priority, Pre-emptive SJF and RR algorithms?

[1.5 marks]

Continued ...

QUESTION 3

a) When using semaphores, a process invokes the `wait()` operation before accessing its critical section, followed by the `signal()` operation upon completion of its critical section.

i) Consider reversing the order of these two operations - first calling `signal()`, then calling `wait()`. What would be the most likely outcome of this?

```
Semaphore S;
signal(S);
...
wait(S);
```

[0.5 + 1 = 1.5 marks]

ii) Suppose a process replaces `signal()` with `wait()` - first calling `wait()`, then again calling `wait()`.What would be the most likely outcome of this?

```
Semaphore S;
wait(S);
...
wait(S);
```

[0.5 + 1 = 1.5 marks]

b) Consider the following system with FIVE processes and THREE resource types:

- A has 10 instances
- B has 10 instances
- C has 10 instances

The snapshot of the system is given as:

Process	Allocation A B C	Max A B C
P0	0 2 2	8 4 4
P1	2 2 2	4 2 2
P2	2 2 2	6 6 2
P3	2 2 2	2 4 2
P4	4 0 0	6 6 8

Available A B C
0 2 2

Answer the following questions using the banker's algorithm:

i) What is the content of the matrix *Need*?

[2.5 marks]

Continued ...

ii) Find a sequence of processes that will get the system into a safe state, if there is any. Show all required steps.

[5 marks]

iii) P4 makes a request for 2 instances of B and 1 instance of C. Check whether the request can be granted.

[1.5 marks]

Continued ...

QUESTION 4

a) One possible implementation of memory protection is the use of a pair of base and limit registers. Discuss how these two registers allow for the protection of memory.

[4 marks]

b) Consider a computer system with a 32-bit logical address and 4 KB page size. The system supports up to 4096 MB of physical memory. How many entries are there in each of the following:

- A conventional single-level page table
- An inverted page table

[1 + 1 = 2 marks]

c) Identify the page numbers and offsets for the following address references (Given as decimal numbers). Assume the page size is 8 KB.

- 18466
- 30000
- 256
- 8264

[0.5 + 0.5 + 0.5 + 0.5 = 2 marks]

d) Standard swapping involves moving processes between main memory and a backing store like a fast disk. Assume the following characteristics:

- Size of user process to be swapped out/to be swapped in: 100 MB
- Transfer rate of backing store (hard disk)/second: 25 MB s^{-1}
- Average Latency Time: 10 ms

Identify:

i) Total transfer time to main memory from backing store/from main memory to backing store (in s).
ii) Swap-out/Swap-in time which includes both transfer time and average latency time (in ms).
iii) Total Swap Time which includes both swap-out and swap-in times (in ms).

[1 + 0.5 + 0.5 = 2 marks]

e) Consider a system that uses a three-level page table with page size of 8KB and 64-bit logical addresses. The first 16 bits of the address serve as the index into the first-level page table, and the last 12 bits serve as the index into the third-level page table.

i) How many bits are needed to specify the second-level index?
ii) How many entries are in a level-one page table?
iii) How many entries are in a level-two page table?
iv) How many pages are in the virtual address space?

[0.5 + 0.5 + 0.5 + 0.5 = 2 marks]

Continued ...

QUESTION 5

a) Paging loads pages only as they are demanded (needed) during program execution. Discuss how a pure demand paging scheme is implemented.

[4 marks]

b) Consider the following page reference string:

3 4 5 3 1 3 4 6 5 2 4 6 2 6 4

Assume the availability of FOUR frames and all the frames are initially empty. Consider the First-In-First Out (FIFO), Least Recently Used (LRU) and Optimal (OPT) page replacement algorithms for Virtual-memory management.

- Identify the number of page faults that would occur if FIFO is used.
- Identify the number of page faults that would occur if LRU is used.
- Identify the number of page faults that would occur if OPT is used.

Show all steps involved.

[2 + 2 + 2 = 6 marks]

c) Given a system with four page frames, the following table indicates page number, load time, last reference time, dirty bit, and reference bit.

Page number	Load time	Last reference time	Dirty bit	Reference bit
0	15	65	1	0
1	20	45	1	1
2	10	80	0	0
3	30	55	0	1

Identify the victim for the following algorithms.

i) First-In-First-Out (FIFO)
ii) Least-Recently-Used (LRU)
iii) Reference bit – Dirty bit combination
iv) Second Chance

[0.5 + 0.5 + 0.5 + 0.5 = 2 marks]

Continued ...

QUESTION 6

a) In order to manage open files, the operating system needs several pieces of information. List and discuss these pieces of information.

[3 marks]

b) For file systems, an allocation method refers to how disk blocks are allocated for files. Explain how contiguous allocation work, and discuss its advantages and disadvantages.

[3 marks]

c) A directory is a collection of nodes containing information to the files that it seeks to contain. Provide a list of data structures that can be used to store this information and discuss the problems that may arise with these data structures.

[4 marks]

d) Consider a file system that uses inodes to represent files. Disk blocks are 8-KB in size and a pointer to a disk block requires 4 bytes. This file system has 12 direct disk blocks, plus single, double, and triple indirect disk blocks. What is the maximum size of a file that can be stored in this file system?

[2 marks]

Continued ...

QUESTION 7

a) Suppose that a disk drive has 2000 cylinders, numbered 0 to 1999. The drive is currently serving a request at cylinder 1145, and the previous request was at cylinder 1125. The queue of pending requests, in FIFO order, is

55, 1450, 815, 1755, 930, 1405, 1105, 1810, 115

Starting from the current head position, what is the schedule and the total distance that the disk arm moves to satisfy all the pending requests, for each of the following disk-scheduling algorithms?

i. FCFS
ii. SCAN

[3 + 3 = 6 marks]

b) Mirroring is a technique employed in RAID schemes to improve reliability. Discuss mirroring.

[2 marks]

c) Linux adopts a three level paging that can be used in both 32-bit and 64-bit architectures. What are the FOUR parts that a linear address in Linux can contain?

[2 marks]

d) Windows implement one-to-one mapping of user threads to kernel threads. What does a thread context contain in Windows?

[2 marks]

End of Exam